PRINCIPLES AND PRACTICE

Augmented REALITY

Dieter SCHMALSTIEG
Tobias HÖLLERER

Augmented Reality

Augmented Reality

Principles and Practice

Dieter Schmalstieg
Tobias Höllerer

Addison-Wesley

Boston • Columbus • Indianapolis • New York • San Francisco • Amsterdam • Cape Town
Dubai • London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City
São Paulo • Sydney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data

Names: Schmalstieg, D. (Dieter), author. | Höllerer, Tobias, 1970– author.
Title: Augmented reality : principles and practice / Dieter Schmalstieg, Tobias Höllerer.
Description: Boston : Addison-Wesley, 2016. | Includes bibliographical references and index.
Identifiers: LCCN 2016009049 | ISBN 9780321883575 (pbk. : alk. paper)
Subject: LCSH: Augmented reality.
Classification: LCC QA76.9.A94 S36 2016 | DDC 006.8—dc23
LC record available at https://lccn.loc.gov/2016009049

ISBN-13: 978-0-321-88357-5
ISBN-10: 0-321-88357-8
Text printed in the United States on recycled paper at RR Donnelley in Crawfordsville, Indiana.
First printing, June 2016

Publisher
Mark L. Taub

Executive Editor
Laura Lewin

Development Editor
Susan Brown Zahn

Managing Editor
Sandra Schroeder

Full-Service Production Manager
Julie B. Nahil

Project Editor
Thistle Hill Publishing Services

Copy Editor
Jill Hobbs

Indexer
John S. Lewis

Proofreader
Anna Popick

Technical Reviewers
Reinhold Behringer
Doug Bowman
Kiyoshi Kiyokawa

Editorial Assistant
Olivia Basegio

Cover Designer
Chuti Prasertsith

Compositor
Shepherd, Inc.

To Ursula, Katharina, and Florian
—Dieter

To Julie, Clara, and Luisa
—Tobias

This page intentionally left blank

Contents

This page intentionally left blank

Preface

Over the past 20 years, the use of information technology has undergone a clear transition from stationary office and desktop computing—first to the web, then to social media, and then to mobile computing. Sales of smartphones and tablet computers have far outpaced the sales of conventional desktop PCs for years now, even if one places laptop or notebook computers within the desktop category.

While the predominant user interface style of today has not radically departed from the *desktop computing* of the 1990s (or the 1981 Xerox Star, for that matter), the way members of the young generation today attain computer literacy has changed: Apps and cloud computing are replacing the computer desktop in many cases. Computing has shifted from office or home office work to an anywhere-and-anytime activity.

Enter: Augmented Reality

As users move away from the desktop, it increasingly makes sense to include the physical world in our computing experience. Given that the physical world is not flat and is not composed of written documents, a new user interface metaphor becomes necessary. **Augmented reality** (AR) has the potential to become the leading user interface metaphor for *situated computing*. Augmented reality has the unique quality of providing a direct link between the physical reality and virtual information about that reality. The world becomes the user interface, leading to the familiar proclamation:

Back to the real world!

Virtual reality, the vision of immersing ourselves in artificial worlds, has propelled the development of game consoles with amazing 3D graphics and led to consumer devices such as head-mounted displays and gesture-tracking devices. Even so, a user interface metaphor such as virtual reality, which by definition monopolizes our attention, is not necessarily a good fit for everyday and spontaneous use of computing.

Instead, we increasingly rely on computer interfaces that make possible casual use and provide information in small, easily understood portions. We feel a need for **ubiquitous computing**. This can take the form of *calm* computing, which operates behind the scenes without the user intervening or even consciously noticing. If ubiquitous *interaction* is required, though, augmented reality excels as an appropriate user interface technology.

Why a Book on Augmented Reality?

Multiple overlapping research fields are contributing to the development of augmented reality, and the associated body of knowledge is growing fast. The authors of this book have been contributing to this body of knowledge as researchers since the 1990s. However, the main motivation for this book came from teaching classes on augmented reality at the authors' home institutions at Graz University of Technology and the University of California, Santa Barbara. In the preparation for these classes, it became evident that no single text is available that covers both the breadth and the depth of this rapidly evolving field. Some notes were available from tutorials at various conferences, several of which the authors were involved in, starting as early as SIGGRAPH 2001. A lot of ground has been covered since then, and the authors were motivated to assemble all this knowledge in a systematic way with an eye toward both novel concepts and practical information. Hence, this book was born.

What's in the Book?

As its title suggests, the book strives for a compromise between principles and practice. Our objective was to make it interesting and usable for both academic researchers and practitioners, especially engineers, who are interested in augmented reality applications. The book, therefore, is intended to be usable both as a textbook and as a reference. To get the most out of it, readers should have a basic understanding of computer science in general, and some knowledge of, and interest in, computer graphics and computer vision is helpful. We don't hesitate to refer to existing literature that explains specific aspects of the necessary background in more detail than we could within the constraints of a single volume. At the same time, we were careful to introduce and clearly explain any specific augmented reality concepts that go beyond basic knowledge, to make the book self-contained. Using the following structure, we present the technical and methodological foundations of AR.

Chapter 1, "Introduction to Augmented Reality," sets the stage by presenting a working definition of augmented reality, providing a brief history of the field, and then walking the reader through various application examples of this powerful real-world user interface technology. We conclude the chapter with a contextualization within the spectrum of related technologies and research fields.

Chapter 2, "Displays," deals with displays, a fundamental enabling technology for augmented reality. Based on some foundations of visual perception, various display technologies that are suitable for augmented reality are discussed—in particular, head-mounted displays, handheld displays, and projective displays. The chapter also discusses nonvisual displays, such as auditory and haptic devices.

Chapter 3, "Tracking," gives an introduction to tracking, one of the core technologies underlying augmented reality. We first discuss the characteristics that are necessary to understand how tracking—and measurement systems in general—work. We then discuss traditional stationary tracking systems and compare them to mobile sensors. Optical tracking as the most prominent tracking technology is given extensive treatment. The chapter concludes by sketching the principles of sensor fusion.

Chapter 4, “Computer Vision for Augmented Reality,” picks up the issue of optical tracking from Chapter 3 and gives a detailed account of computer vision algorithms for real-time pose estimation, i.e. for determining a camera's viewing position and orientation from the observed imagery. To make this topic manageable and address readers with a wide variety of backgrounds, the chapter is structured along a suite of case studies. Every case study introduces only the knowledge necessary for it to be self-contained, so the reader does not have to accumulate in-depth knowledge of computer vision first. Moreover, advanced mathematical topics, which in practice are often used as a black box by relying on a software library such as OpenCV, are marked so that the reader can safely skip over them.

Chapter 5, “Calibration and Registration,” deals with methods for calibration and registration of the devices used in augmented reality. Calibration of the digital cameras used for the optical tracking described in Chapter 3 is a necessary prerequisite to deliver repeatable, accurate behavior in augmented reality applications. Registration is the process that aligns the physical and virtual parts of the augmented reality experience geometrically, thereby giving rise to the illusion of a coherent mixed environment.

Chapter 6, “Visual Coherence,” focuses on a family of computer graphics techniques that together produce a seamlessly blended view of real and virtual objects. It includes phenomena such as correct occlusion between virtual and real objects, or correct shadowing between virtual and real objects. We also explain diminished reality, or the removal of real objects from a scene, and examine the simulation of physical cameras.

Chapter 7, “Situated Visualization,” is dedicated to visualization techniques. Visualization has the objective of making information comprehensible. In the context of augmented reality, this means that the computer-generated information that is geometrically registered to objects in the physical scene must be positioned and styled in such a way that it can be easily understood by its users. We deal with both two-dimensional augmentations (such as textual labels) and three-dimensional augmentations (such as synthesized views of the interior of objects, so-called “ghostings”).

Chapter 8, “Interaction,” examines the various interaction techniques and interaction styles that are relevant for augmented reality applications. The topics range from simple situated information browsing to full three-dimensional interaction. We specifically discuss props, widgets, and hand-based interaction, and the connection of augmented reality to tangible user interfaces of various forms. We also take a look at multimodal and agent-based interfaces for augmented reality.

Chapter 9, “Modeling and Annotation,” is concerned with the topic of interactive modeling—that is, the creation of new geometric content through augmented reality. User interfaces that are embedded in a three-dimensional environment provide a powerful approach for re-creating a digital version of this environment. This capability is invaluable for all applications that deal with visual computing.

Chapter 10, “Authoring,” discusses authoring approaches for augmented reality. The content of augmented reality presentations and information databases needs to be designed and created the same way that web content is authored today. Augmented reality content can be authored with conventional tools or in augmented reality itself. Authoring is concerned with aspects of the application that go beyond geometric and visual properties—in particular, establishing the semantics and the behavior of the application. Preferably, authoring should be content-driven and require no or only minimal traditional programming effort. We discuss various approaches to address this need, and also examine recent efforts to combine augmented reality authoring with emerging open web standards.

Chapter 11, “Navigation,” deals with navigational guidance—a particularly relevant aspect of augmented reality as a user interface. Orientation in unfamiliar environments is an important application challenge involving mobile information systems. We present an overview of navigational guidance techniques implemented using augmented reality, and compare them to digital maps.

Chapter 12, “Collaboration,” investigates collaboration. Augmented reality has strong potential as a medium that can be used for communication among individuals. This encompasses both co-located collaboration, which is enriched by the additional cues afforded by a shared augmented reality system, and remote collaboration, which can be significantly supported by augmented reality technology and, in the process, provide new forms of remote presence.

Chapter 13, “Software Architectures,” analyzes the underlying architectures of augmented reality systems. Augmented reality has complex requirements, as it must combine aspects of real-time systems, multimedia systems, and often also distributed systems. Combining these requirements in a flexible way that can be mastered by an application programmer is a difficult endeavor. We discuss various architectural patterns such as distributed objects, dataflow systems, and scene graphs, and present a number of case studies.

Chapter 14, “ The Future,” reviews possible trajectories of augmented reality as it moves from a research field with demonstrated usefulness in prototype applications to potentially universal consumer adoption. As part of this effort, the chapter considers which roadblocks and unresolved issues remain to be overcome. It also summarizes trends and insights from all of the material presented in this book and sketches a future research agenda.

How to Use the Book and the Related Material

How you use this book will depend on your relationship to the field of augmented reality and the degree and focus of your interest. We discuss three types of roles that this relationship or interest might take.

If you are a developer: Professional developers can use the book for inspiration and guidance in the design, implementation, and evaluation of augmented reality applications. Readers with such backgrounds will find useful information on hardware setups in the display, tracking, and interaction chapters. They will benefit from the chapters on visual coherence, visualization, and authoring for the development of application content, and learn about appropriate registration technologies in the tracking, computer vision, and calibration chapters. User interface design is informed by the chapter on interaction and following chapters. Finally, the chapter on software architectures provides important information for actual implementation work.

If you are a teacher: The book is useful as a text for several different types of university-level courses. A graduate course on augmented reality can use it as the primary textbook. A course on computer graphics or visual computing could use the chapters on visual coherence and visualization as an introduction to graphical aspects of augmented reality. A course on computer vision can use the chapters on tracking and registration for teaching important real-time computer vision techniques. A human–computer interaction course can utilize the chapters on interaction, modeling, authoring, navigation, and collaboration to provide detailed coverage of augmented reality concepts.

If you are a researcher: This book can serve as a comprehensive reference guide for researchers interested in the development or evaluation of experimental augmented reality applications. The research agenda in the concluding chapter also provides researchers and students with a list of important questions to be addressed in the field.

Companion Website

The companion website to the book can be found at the following address:

http://www.augmentedrealitybook.org

Augmented reality is rapidly evolving. To make this book a dynamic working document, this companion website provides additional information, including teaching materials. This site contains information and links related to the latest augmented reality research and applications. This is an open effort, so readers are invited to contribute to this collection. Your comments will help us to update the website, as well as future editions of this book.

Register your copy of *Augmented Reality* at informit.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN (9780321883575) and click Submit. Once the process is complete, you will find any available bonus content under "Registered Products."

Acknowledgments

This book would not have been possible without the encouragement and expertise of many friends and colleagues. First, we offer our gratitude to the reviewers who provided invaluable insights and suggestions for improvements: Reinhold Behringer, Doug Bowman, André Ferko, Steffen Gauglitz, Kiyoshi Kiyokawa, Tobias Langlotz, Vincent Lepetit, Gerhard Reitmayr, Chris Sweeney, and Daniel Wagner.

Second, we want to thank our editors at Addison-Wesley: Peter Gordon, who believed in the idea of this book and helped us establish the publishing contract, and Laura Lewin and Olivia Basegio, who continuously encouraged us and provided us with great advice.

Third, we want to thank all colleagues that provided us with additional image materials: Aaron Stafford, Alessandro Mulloni, Alexander Plopski, Andreas Butz, Andreas Geiger, Andreas Hartl, Andrei State, Andrew Maimone, Andy Gstoll, Ann Morrison, Anton Fuhrmann, Anton van den Hengel, Arindam Day, Blair MacIntyre, Brigitte Ludwig, Bruce Thomas, Christian Pirchheim, Christian Reinbacher, Christian Sandor, Claudio Pinhanez, Clemens Arth, Daniel Wagner, David Mizell, Denis Kalkofen, Domagoj Baričević, Doreé Seligmann, Eduardo Veas, Erick Mendez, Ernst Kruijff, Ethan Eade, Florian Ledermann, Gerd Hesina, Gerhard Reitmayr, Gerhard Schall, Greg Welch, Gudrun Klinker, Hannes Kaufmann, Henry Fuchs, Hiroyuki Yamamoto, Hrvoje Benko, István Barakonyi, Ivan Sutherland, Jan Herling, Jens Grubert, Jonathan Ventura, Joseph Newman, Julien Pilet, Kiyoshi Kiyokawa, Lukas Gruber, Mark Billinghurst, Markus Oberweger, Markus Tatzgern, Martin Hirzer, Matt Swoboda, Matthias Straka, Michael Gervautz, Michael Kenzel, Michael Marner, Morten Fjeld, Nassir Navab, Oliver Bimber, Pascal Fua, Pascal Lagger, Peter Kán, Peter Mohr, Peter Weir, Philipp Descovic, Qi Pan, Ralph Schönfelder, Raphael Grasset, Remo Ziegler, Simon Julier, Stefan Hauswiesner, Stefanie Zollmann, Steffen Gauglitz, Steve Feiner, Taehee Lee, Takuji Narumi, Thanh Nguyen, Thomas Richter-Trummer, Tom Drummond, Ulrich Eck, Vincent Lepetit, Wayne Piekarski, William Steptoe, Wolfgang Broll, and Zsolt Szalavári.

Fourth, we want to thank all colleagues and students at Graz University of Technology and University of California, Santa Barbara, for the countless discussions on and off the topics in this book and for providing great working environments.

Finally, we would like to thank our families for supporting us and being patient during the not-so-short time of creating this book!

Dieter Schmalstieg
Graz, Austria, April 2016

Tobias Höllerer
Santa Barbara, California, April 2016

This page intentionally left blank

About the Authors

Dieter Schmalstieg is full professor and head of the Institute of Computer Graphics and Vision at Graz University of Technology (TUG), Austria. His current research interests are augmented reality, virtual reality, real-time graphics, user interfaces, and visualization. He received Dipl.-Ing. (1993), Dr.techn. (1997), and Habilitation (2001) degrees from Vienna University of Technology. Dr. Schmalstieg is author or coauthor of more than 300 peer-reviewed scientific publications. His organizational roles include associate editor in chief of *IEEE Transactions on Visualization and Computer Graphics*, member of the editorial advisory board of *Computers & Graphics* and of Springer's *Virtual Reality* journal, member of the steering committee of the IEEE International Symposium on Mixed and Augmented Reality, chair of the EUROGRAPHICS working group on Virtual Environments (1999–2010), key researcher of the K-Plus Competence Center for Virtual Reality and Visualization in Vienna, and key researcher of the Know-Center in Graz. In 2002, Dr. Schmalstieg received the START Career Award presented by the Austrian Science Fund. In 2012, he received the IEEE Virtual Reality Technical Achievement Award for seminal contributions to the field of augmented reality. He was elected as a senior member of IEEE, as a member of the Austrian Academy of Sciences, and as a member of the Academia Europaea. Since 2008, he has also been director of the Christian Doppler Laboratory for Handheld Augmented Reality.

Tobias Höllerer is professor of computer science at the University of California, Santa Barbara, where he co-directs the Four Eyes Laboratory, conducting research in the "four I's" of Imaging, Interaction, and Innovative Interfaces. Dr. Höllerer holds a Diploma in Informatics from the Technical University of Berlin, as well as an M.S., an M.Phil., and a Ph.D. in computer science from Columbia University. He is a recipient of the U.S. National Science Foundation's CAREER Award for his work on "Anywhere Augmentation," which enables mobile computer users to place annotations in 3D space wherever they go. He is an IEEE senior member and was named an ACM Distinguished Scientist in 2013. Dr. Höllerer is author or coauthor of more than 150 peer-reviewed journal and conference publications in the areas of augmented and virtual reality, information visualization, 3D displays and interaction, mobile and wearable computing, and social computing. Several of these publications have been selected for Best Paper or Honorable Mention Awards at such venues as the IEEE International Symposium on Mixed and Augmented Reality (ISMAR), IEEE Virtual Reality (VR), ACM Virtual Reality Software and Technology, ACM User Interface Software and Technology, ACM MobileHCI, IEEE SocialCom, and IEEE CogSIMA. Dr. Höllerer is an associate editor of *IEEE Transactions on Visualization and Computer Graphics*. Among his many organizational roles for scientific conferences, he served as program chair for IEEE VR 2015, ICAT 2013, and IEEE ISMAR 2009 and 2010; as general chair of IEEE ISMAR 2006; and as member of the steering committee of IEEE ISMAR.

This page intentionally left blank

CHAPTER 1

INTRODUCTION TO AUGMENTED REALITY

Virtual reality is becoming increasingly popular, as computer graphics have progressed to a point where the images are often indistinguishable from the real world. However, the computer-generated images presented in games, movies, and other media are detached from our physical surroundings. This is both a virtue—everything becomes possible—and a limitation.

The limitation comes from the main interest we have in our daily life, which is not directed toward some virtual world, but rather toward the *real world* surrounding us. Smartphones and other mobile devices provide access to a vast amount of information, anytime and anywhere. However, this information is generally disconnected from the real world. Consumers with an interest in retrieving online information from and about the real world, or linking up online information with the real world, must do so individually and indirectly, which, in turn, requires constant cognitive effort.

In many ways, enhancing mobile computing so that the association with the real world happens automatically seems an attractive proposition. A few examples readily illustrate this idea's appeal. Location-based services can provide personal navigation based on the Global Positioning System (GPS), while barcode scanners can help identify books in a library or products in a supermarket. These approaches require explicit actions by the user, however, and are rather coarse grained. Barcodes are useful for identifying books, but not for naming mountain peaks during a hiking trip; likewise, they cannot help in identifying tiny parts of a watch being repaired, let alone anatomic structures during surgery.

Augmented reality holds the promise of creating direct, automatic, and actionable links between the physical world and electronic information. It provides a simple and immediate user interface to an electronically enhanced physical world. The immense potential of augmented reality as a paradigm-shifting user interface metaphor becomes apparent when we review the most recent few milestones in human–computer interaction: the emergence of the World Wide Web, the social web, and the mobile device revolution.

The trajectory of this series of milestones is clear: First, there was an immense increase in access to online information, leading to a massive audience of information consumers. These consumers were subsequently enabled to also act as information producers and communicate with one another, and finally were given the means to manage their communications from anywhere, in any situation. Yet, the physical world, in which all this information retrieval, authoring, and communication takes place, was not readily linked to the users' electronic activity. That is, the model was stuck in a world of abstract web pages and services without directly involving the physical world. A lot of technological advancement has occurred in the field of location-based computing and services, which is sometimes referred to as *situated computing*. Even so, the user interfaces to location-based services remain predominantly rooted in desktop-, app-, and web-based usage paradigms.

Augmented reality can change this situation, and, in doing so, redefine information browsing and authoring. This user interface metaphor and its enabling technologies form one of today's most fascinating and future-oriented areas of computer science and application development. Augmented reality can overlay computer-generated information on views of the real world, amplifying human perception and cognition in remarkable new ways.

After providing a working definition of augmented reality, we will briefly review important developments in the history of the research field, and then present examples from various application areas, showcasing the power of this physical user interface metaphor.

Definition and Scope

Whereas virtual reality (VR) places a user inside a completely computer-generated environment, augmented reality (AR) aims to present information that is directly registered to the physical environment. AR goes beyond mobile computing in that it bridges the gap between virtual world and real world, both spatially and cognitively. With AR, the digital information appears to become part of the real world, at least in the user's perception.

Achieving this connection is a grand goal—one that draws upon knowledge from many areas of computer science, yet can lead to misconceptions about what AR really is. For example, many people associate the visual combination of virtual and real elements with the special effects in movies such as *Jurassic Park* and *Avatar*. While the computer graphics techniques used in movies may be applicable to AR as well, movies lack one crucial aspect of AR—interactivity. To avoid such confusion, we need to set a scope for the topics discussed in this book. In other words, we need to answer a key question: What is AR?

The most widely accepted definition of AR was proposed by Azuma in his 1997 survey paper. According to Azuma [1997], AR must have the following three characteristics:

- Combines real and virtual
- Interactive in real time
- Registered in 3D

This definition *does not require* a specific output device, such as a head-mounted display (HMD), nor does it limit AR to visual media. Audio, haptics, and even olfactory or gustatory AR are included in its scope, even though they may be difficult to realize. Note that the definition does require real-time *control* and spatial *registration*, meaning precise real-time alignment of corresponding virtual and real information. This mandate implies that the user of an AR display can at least exercise some sort of interactive viewpoint control, and the computer-generated augmentations in the display will remain registered to the referenced objects in the environment.

While opinions on what qualifies as real-time performance may vary depending on the individual and on the task or application, interactivity implies that the human–computer interface operates in a tightly coupled feedback loop. The user continuously navigates the AR scene and controls the AR experience. The system, in turn, picks up the user's input by tracking the user's viewpoint or pose. It registers the pose in the real world with the virtual content, and then presents to the user a *situated visualization* (a visualization that is registered to objects in the real world).

We can see that a complete AR system requires at least three components: a tracking component, a registration component, and a visualization component. A fourth component—a spatial model (i.e., a database)—stores information about the real world and about the virtual world (Figure 1.1). The real-world model is required to serve as a reference for the tracking

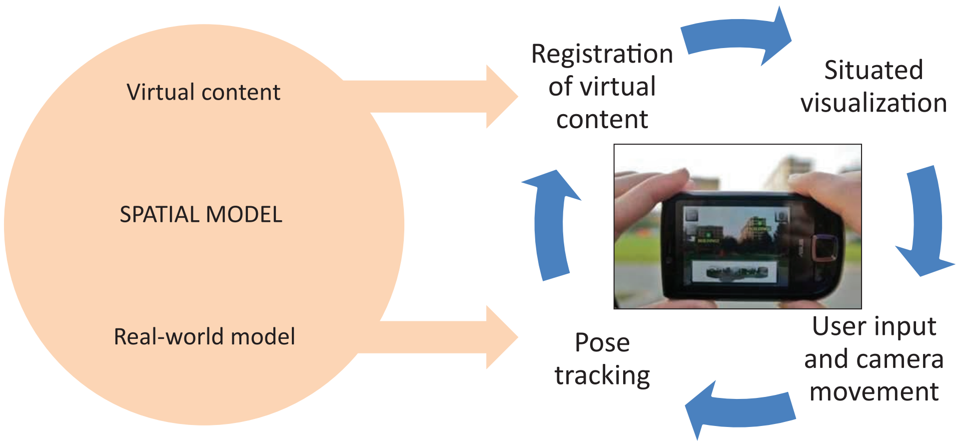

Figure 1.1 AR uses a feedback loop between human user and computer system. The user observes the AR display and controls the viewpoint. The system tracks the user's viewpoint, registers the pose in the real world with the virtual content, and presents situated visualizations.

component, which must determine the user's location in the real world. The virtual-world model consists of the content used for the augmentation. Both parts of the spatial model must be registered in the same coordinate system.

A Brief History of Augmented Reality

While one could easily go further back in time to find examples in which informational overlays were layered on top of the physical world, suffice it to say that the first annotations of the physical world with *computer-generated* information occurred in the 1960s. Ivan Sutherland can be credited with starting the field that would eventually turn into both VR and AR. In 1965, he postulated the *ultimate display* in an essay that contains the following famous quote:

> The ultimate display would, of course, be a room within which the computer can control the existence of matter. A chair displayed in such a room would be good enough to sit in. Handcuffs displayed in such a room would be confining, and a bullet displayed in such a room would be fatal. With appropriate programming such a display could literally be the Wonderland into which Alice walked.

Sutherland's [1965] essay includes more than just an early description of immersive displays, however. It also contains a quote that is less often discussed, but that clearly anticipates AR:

> The user of one of today's visual displays can easily make solid objects transparent—he can "see through matter!"

Figure 1.2 The Sword of Damocles was the nickname of the world's first head-mounted display, built in 1968. Courtesy of Ivan Sutherland.

Shortly thereafter, Sutherland constructed the first VR system. In 1968, he finished the first head-mounted display [Sutherland 1968]. Because of its weight, it had to be suspended from the ceiling and was appropriately nicknamed "Sword of Damocles" (Figure 1.2). This display already included head tracking and used see-through optics.

Advances in computing performance of the 1980s and early 1990s were ultimately required for AR to emerge as an independent field of research. Throughout the 1970s and 1980s, Myron Krueger, Dan Sandin, Scott Fisher, and others had experimented with many concepts of mixing human interaction with computer-generated overlays on video for interactive art experiences. Krueger [1991], in particular, demonstrated collaborative interactive overlays of graphical annotations among participant silhouettes in his Videoplace installations around 1974.

The year 1992 marked the birth of the term "augmented reality." This term first appeared in the work of Caudell and Mizell [1992] at Boeing, which sought to assist workers in an airplane factory by displaying wire bundle assembly schematics in a see-through HMD (Figure 1.3).

Figure 1.3 Researchers at Boing used a see-through HMD to guide the assembly of wire bundles for aircraft. Courtesy of David Mizell.

In 1993, Feiner et al. [1993a] introduced KARMA, a system that incorporated knowledge-based AR. This system was capable of automatically inferring appropriate instruction sequences for repair and maintenance procedures (Figure 1.4).

Also in 1993, Fitzmaurice created the first handheld spatially aware display, which served as a precursor to handheld AR. The Chameleon consisted of a tethered handheld liquid-crystal display (LCD) screen. The screen showed the video output of an SGI graphics workstation of the time and was spatially tracked using a magnetic tracking device. This system was capable of showing contextual information as the user moved the device around—for example, giving detailed information about a location on a wall-mounted map.

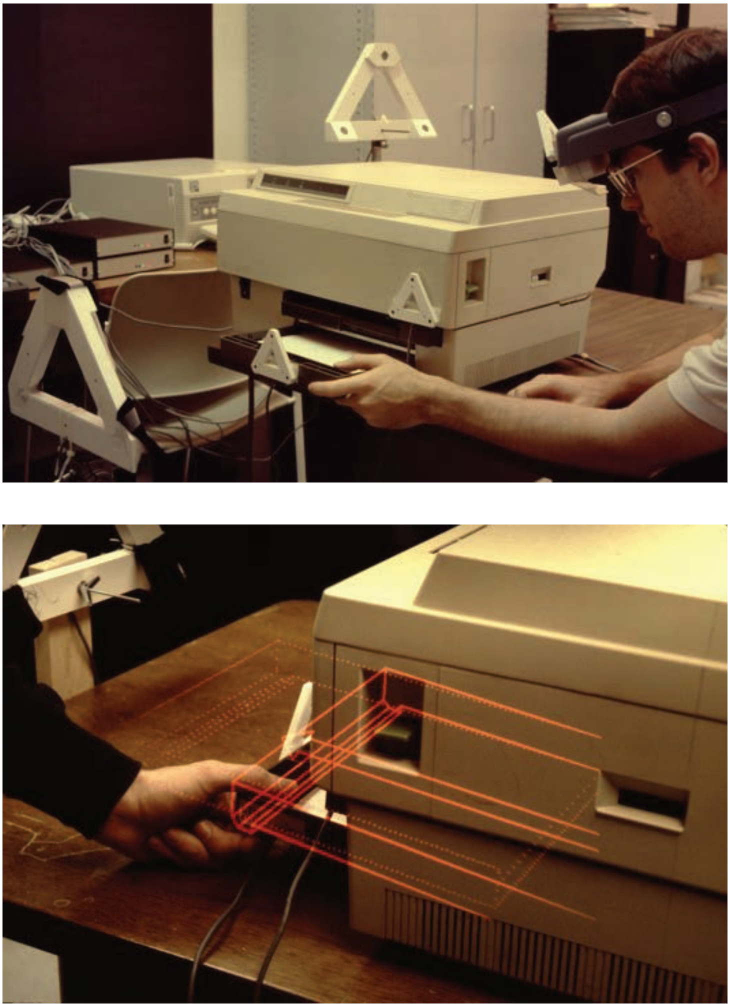

Figure 1.4 (top) KARMA was the first knowledge-driven AR application. (bottom) A user with an HMD could see instructions on printer maintenance. Courtesy of Steve Feiner, Blair MacIntyre, and Doreé Seligmann, Columbia University.

In 1994, State et al. at the University of North Carolina at Chapel Hill presented a compelling medical AR application, capable of letting a physician observe a fetus directly within a pregnant patient (Figure 1.5). Even though the accurate registration of computer graphics on top of a deformable object such as a human body remains a challenge today, this seminal work hints at the power of AR for medicine and other delicate tasks.

Figure 1.5 View inside the womb of an expecting mother. Courtesy of Andrei State, UNC Chapel Hill.

Around the mid-1990s, Steve Mann at the MIT Media Lab implemented, and experimented with, a "reality mediator"—a waist-bag computer with a video see-through HMD (a modified VR4 by Virtual Research Systems) that enabled the user to augment, alter, or diminish visual reality. Through the WearCam project, Mann [1997] explored wearable computing and mediated reality. His work ultimately helped establish the academic field of wearable computing, which, in those early days, had a lot of synergy with AR [Starner et al. 1997].

In 1995, Rekimoto and Nagao created the first true—albeit tethered—handheld AR display. Their NaviCam was connected to a workstation, but was outfitted with a forward-facing camera. From the video feed, it could detect color-coded markers in the camera image and display information on a video see-through view.

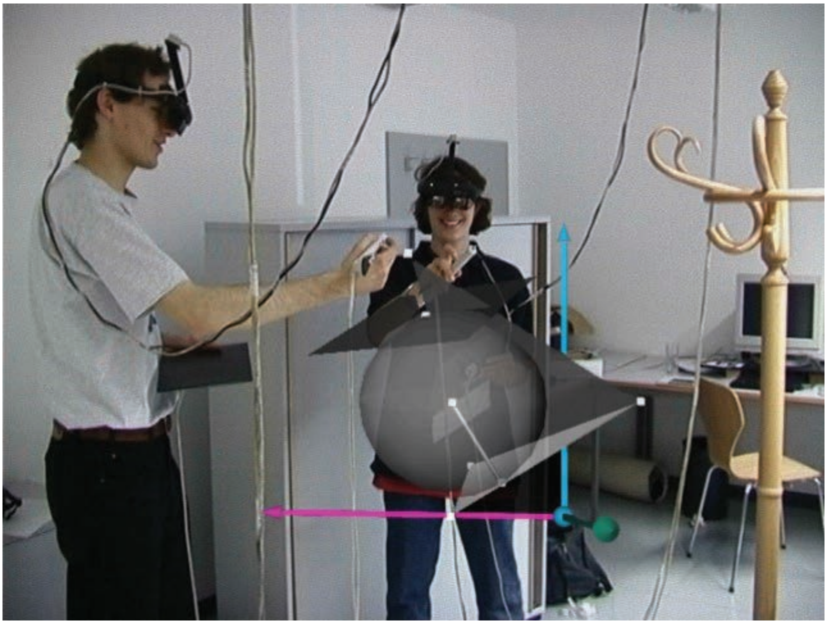

Figure 1.6 One of the applications of the Studierstube system was teaching geometry in AR to high school students. Courtesy of Hannes Kaufmann.

In 1996, Schmalstieg et al. developed Studierstube, the first collaborative AR system. With this system, multiple users could experience virtual objects in the same shared space. Each user had a tracked HMD and could see perspectively correct stereoscopic images from an individual viewpoint. Unlike in multi-user VR, natural communication cues, such as voice, body posture, and gestures, were not affected in Studierstube, because the virtual content was added to a conventional collaborative situation in a minimally obtrusive way. One of the showcase applications was a geometry course [Kaufmann and Schmalstieg 2003], which was successfully tested with actual high school students (Figure 1.6).

From 1997 to 2001, the Japanese government and Canon Inc. jointly funded the Mixed Reality Systems Laboratory as a temporary research company. This joint venture was the largest industrial research facility for mixed reality (MR) research up to that point [Tamura 2000] [Tamura et al. 2001]. Among its most notable achievements was the design of the first coaxial stereo video see-through HMD, the COASTAR. Many of the activities undertaken in the lab were also directed toward the digital entertainment market (Figure 1.7), which plays a very prominent role in Japan.

In 1997, Feiner et al. developed the first outdoor AR system, the Touring Machine (Figure 1.8), at Columbia University. The Touring Machine uses a see-through HMD with GPS and orientation tracking. Delivering mobile 3D graphics via this system required a backpack holding a

Figure 1.7 *RV-Border Guards* was a multiuser shooting game developed in Canon's Mixed Reality Systems Laboratory. Courtesy of Hiroyuki Yamamoto.

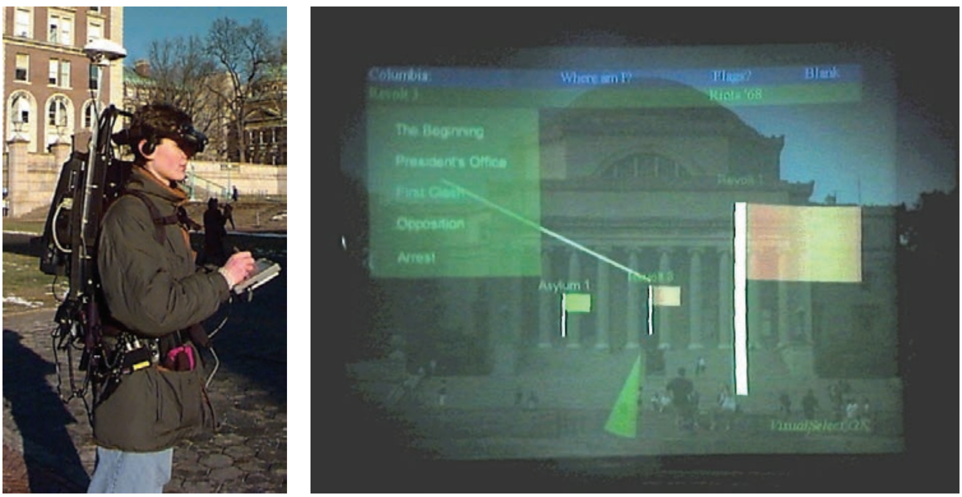

Figure 1.8 The Touring Machine was the first outdoor AR system (left). Image of the *Situated Documentaries* AR campus tour guide running on a 1999 version of the Touring Machine (right). Courtesy of Columbia University.

computer, various sensors, and an early tablet computer for input [Feiner et al. 1997] [Höllerer et al. 1999b].

Just one year later, in 1998, Thomas et al. published their work on the construction of an outdoor AR navigation system, Map-in-the-Hat. Its successor, Tinmith (few people know that this name is actually an acronym for “This is not map in the hat”), evolved into a well-known experimental platform for outdoor AR. This platform was used for advanced applications, such as 3D surveying, but is most famous for delivering the first outdoor AR game, *ARQuake* (Figure 1.9). This game, which is a port of the popular first-person shooter application *Quake* to Tinmith, places the user in the midst of a zombie attack in a real parking lot.

In the same year, Raskar et al. [1998] at the University of North Carolina at Chapel Hill presented the Office of the Future, a telepresence system built around the idea of structured light-scanning and projector-camera systems. Although the required hardware was not truly practical for everyday use at the time, related technologies, such as depth sensors and camera-projection coupling, play a prominent role in AR and other fields today.

Until 1999, no AR software was available outside specialized research labs. This situation changed when Kato and Billinghurst [1999] released ARToolKit, the first open-source software

Figure 1.9 Screenshot of ARQuake, the first outdoor AR game. Courtesy of Bruce Thomas and Wayne Piekarski.

Figure 1.10 A person holding a square marker of ARToolKit, the popular open-source software framework for AR. Courtesy of Mark Billinghurst.

platform for AR. It featured a 3D tracking library using black-and-white fiducials, which could easily be manufactured on a laser printer (Figure 1.10). The clever software design, in combination with the increased availability of webcams, made ARToolKit widely popular.

In the same year, Germany's Federal Ministry for Education and Research initiated a €21 million program for industrial AR, called ARVIKA (Augmented Reality for Development, Production, and Servicing). More than 20 research groups from industry and academia worked on developing advanced AR systems for industrial application, in particular in the German automotive industry. This program raised the worldwide awareness of AR in professional communities and was followed by several similar programs designed to enhance industrial application of the technology.

Another noteworthy idea also appeared in the late 1990s: IBM researcher Spohrer [1999] published an essay on Worldboard, a scalable networked infrastructure for hyperlinked spatially registered information, which Spohrer had first proposed while he was working with Apple's Advanced Technology Group. This work can be seen as the first concept for an AR browser.

After 2000, cellular phones and mobile computing began evolving rapidly. In 2003, Wagner and Schmalstieg presented the first handheld AR system running autonomously on a "personal digital assistant"—a precursor to today's smartphones. One year later, the *Invisible Train* [Pintaric et al. 2005], a multiplayer handheld AR game (Figure 1.11), was experienced by thousands of visitors at the SIGGRAPH Emerging Technologies show floor.

It took several years, until 2008, for the first truly usable natural feature tracking system for smartphones to be introduced [Wagner et al. 2008b]. This work became the ancestor of the

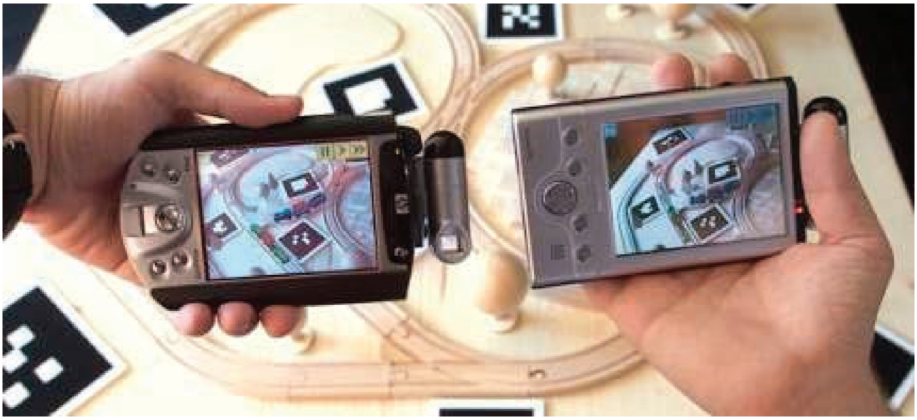

Figure 1.11 The Invisible Train was a handheld AR game featuring virtual trains on real wooden tracks. Courtesy of Daniel Wagner.

popular Vuforia toolkit for AR developers. Other noteworthy achievements in recent years in the area of tracking include the parallel tracking and mapping (PTAM) system of Klein and Murray [2007], which can track without preparation in unknown environments, and the KinectFusion system developed by Newcombe et al. [2011a], which builds detailed 3D models from an inexpensive depth sensor. Today, AR developers can choose among many software platforms, but these model systems continue to represent important directions for researchers.

Examples

In this section, we continue our exploration of AR by examining a set of examples, which showcase both AR technology and applications of that technology. We begin with application domains in which AR technologies demonstrated early success—namely, industry and construction. These examples are followed by applications in maintenance and training, and in the medical domain. We then discuss examples that focus on individuals on the move: personal information display and navigational support. Finally, we present examples illustrating how large audiences can be supported by AR using enhanced media channels in, for example, television, online commerce, and gaming.

Industry and Construction

As mentioned in our brief historic overview of AR, some of the first actual applications motivating the use of AR were industrial in nature, such as Boeing's wire bundle assembly needs and early maintenance and repair examples.

Industrial facilities are becoming increasingly complex, which profoundly affects their planning and operation. Architectural structures, infrastructure, and machines are planned using computer-aided design (CAD) software, but typically many alterations are made during actual construction and installation. These alterations usually do not find their way back into the CAD models. In addition, there may be a large body of legacy structures predating the introduction of CAD for planning as well as the need for frequent changes of the installations—for example, when a factory is adapted for the manufacturing of a new product. Planners would like to compare the "as planned" to the "as is" state of a facility and identify any critical deviations. They would also like to obtain a current model of the facility, which can be used for planning, refurbishing or logistics procedures.

Traditionally, this is done with 3D scanners and off-site data integration and comparison. This process is lengthy and tedious, however, and it results in low-level models consisting of point clouds. AR offers the opportunity to perform on-site inspection, bringing the CAD model to the facility rather than the reverse. Georgel et al. [2007], for example, have developed a technique for still-frame AR that extracts the camera pose from perspective cues in a single image and overlays registered, transparently rendered CAD models (Figure 1.12).

Schönfelder and Schmalstieg [2008] have proposed a system based on the Planar (Figure 1.13), an AR display on wheels with external tracking. It provides fully interactive, real-time discrepancy checking for industrial facilities.

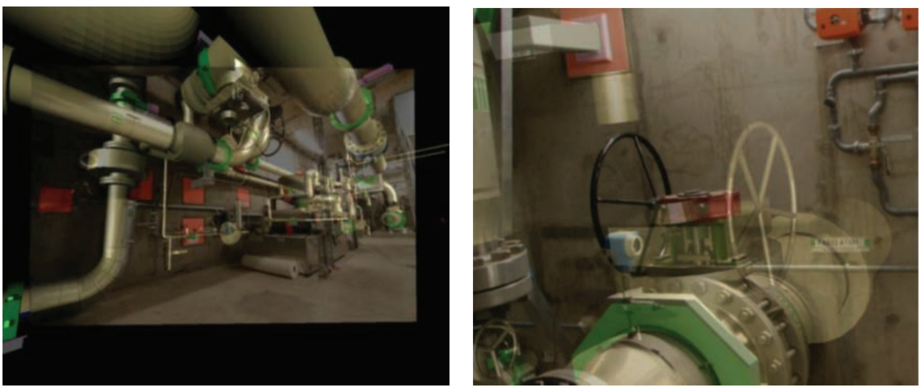

Figure 1.12 AR can be used for discrepancy analysis in industrial facilities. These images show still frames overlaid with CAD information. Note how the valve on the right-hand side was mounted on the left side rather than on the right side as in the model. Courtesy of Nassir Navab.

Figure 1.13 The Planar is a touchscreen display on wheels (left), which can be used for discrepancy analysis directly on the factory floor (right). Courtesy of Ralph Schönfelder.

Utility companies rely on geographic information systems (GIS) for managing underground infrastructure, such as telecommunication lines or gas pipes. The precise locations of the underground assets are required in a variety of situations. For example, construction managers are legally obliged to obtain information on underground infrastructure, so that they can avoid any damage to these structures during excavations. Likewise, locating the reason for outages or updating outdated GIS information frequently requires on-site inspection. In all these cases, presenting an AR view that is derived from the GIS and directly registered on the target site can significantly improve the precision and speed of outdoor work [Schall et al. 2008]. Figure 1.14 shows Vidente, one such outdoor AR visualization system.

Camera-bearing micro-aerial vehicles (drones) are increasingly being used for airborne inspection and reconstructions of construction sites. These drones may have some degree of autonomous flight control, but always require a human operator. AR can be extremely useful in locating the drone (Figure 1.15), monitoring its flight parameters such as position over ground, height, or speed, and alerting the operator to potential collisions [Zollmann et al. 2014].

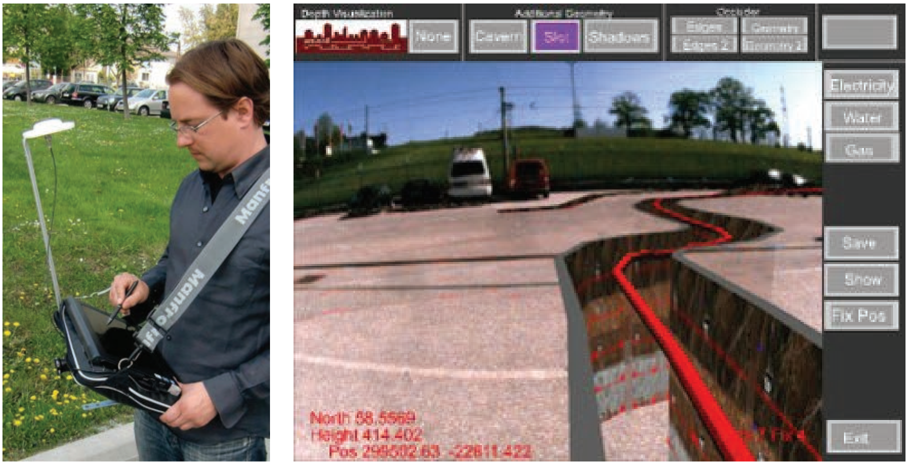

Figure 1.14 Tablet computer with differential GPS system for outdoor AR (left). Geo-registered view of a virtual excavation revealing a gas pipe (right). Courtesy of Gerhard Schall.

Figure 1.15 While the drone has flown far away and is barely visible, its position can be visualized using a spherical AR overlay. Courtesy of Stefanie Zollmann.

Maintenance and Training

Understanding how things work, and learning how to assemble, disassemble, or repair them, is an important challenge in many professions. Maintenance engineers often devote a large amount of time to studying manuals and documentation, since it is often impossible to memorize all procedures in detail. AR, however, can present instructions directly superimposed in the field of view of the worker. This can provide more effective training, but, more importantly, allows personnel with less training to correctly perform the work. Figure 1.16 reveals how AR can assist with the removal of the brewing unit of an automatic coffee maker, and Figure 1.17 shows the disassemble sequence for a valve [Mohr et al. 2015].

If human support is sought, AR can provide a shared visual space for live mobile remote collaboration on physical tasks [Gauglitz et al. 2014a]. With this approach, a remote expert can explore the scene independently of the local user's current camera position and can communicate via

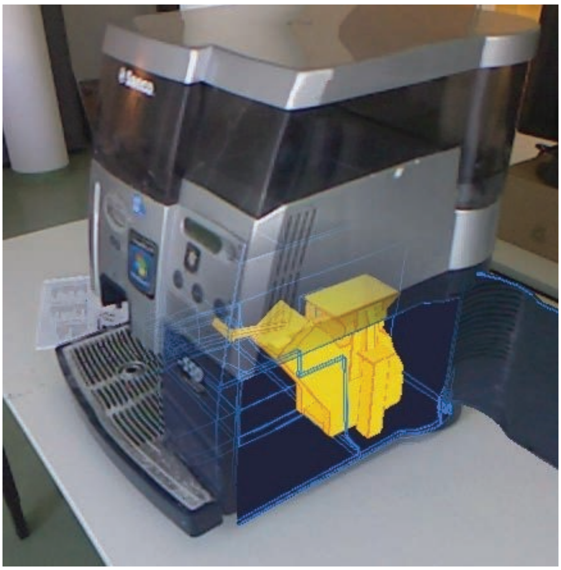

Figure 1.16 Ghost visualization revealing the interior of a coffee machine to guide end-user maintenance. Courtesy of Peter Mohr.

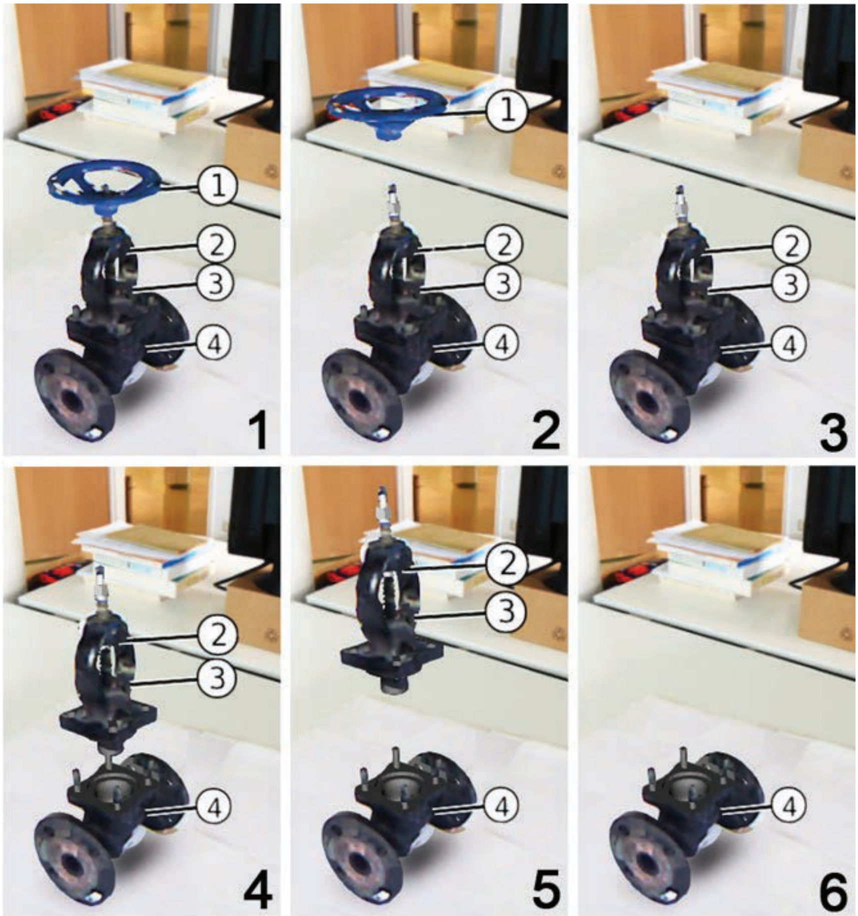

Figure 1.17 Automatically generated disassembly sequence of a valve. Courtesy of Peter Mohr.

spatial annotations that are immediately visible to the local user in the AR view (Figure 1.18). This can be achieved with real-time visual tracking and reconstruction, eliminating the need for preparation or instrumentation of the environment. AR telepresence combines the benefits of live video conferencing and remote scene exploration into a natural collaborative interface.

Medical

The use of X-ray imaging revolutionized diagnostics by allowing physicians to see inside a patient without performing surgery. However, conventional X-ray and computed tomography devices separate the interior view from the exterior view of the patient. AR integrates these

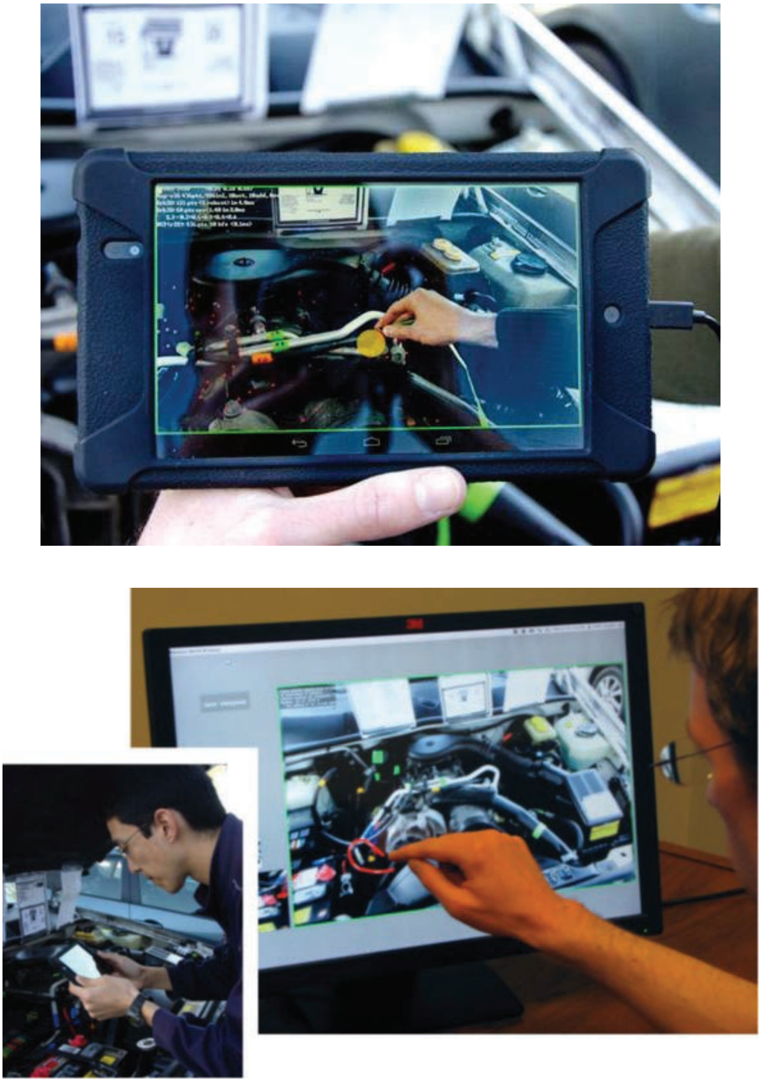

Figure 1.18 A car repair scenario assisted by a remote expert via AR telepresence on a tablet computer (top). The remote expert can draw hints directly on the 3D model of the car that is incrementally transmitted from the repair site (bottom). Courtesy of Steffen Gauglitz.

views, enabling the physician to see directly inside the patient. One such example, which is now commercially available, is the Camera Augmented Mobile C-arm, or CamC (Figure 1.19). A mobile C-arm is used to provide X-ray views in the operating theater. CamC extends these views with a conventional video camera, which is arranged coaxially with the X-ray optics to deliver precisely registered image pairs [Navab et al. 2010]. The physician can transition and blend

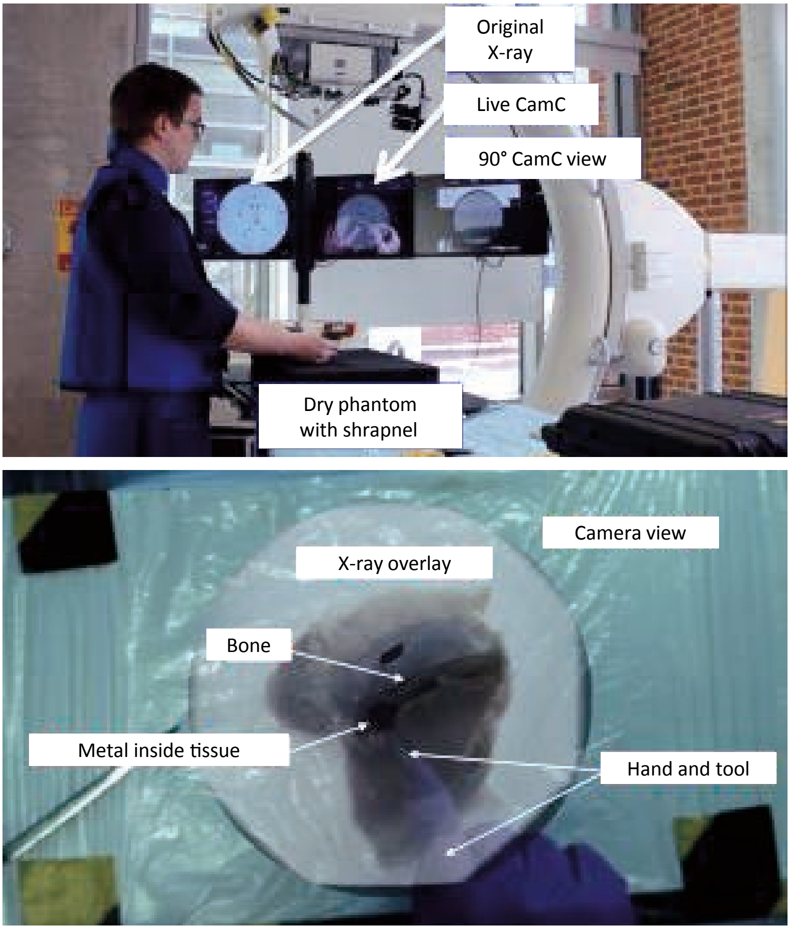

Figure 1.19 The CamC is a mobile C-arm, which allows a physician to seamlessly blend between a conventional camera view and X-ray images. Courtesy of Nassir Navab.

between the inside and outside views as desired. CamC has many clinical applications, including guiding needle biopsies and facilitating orthopedic screw placement.

Personal Information Display

As we have seen, several specific application domains can profit from the use of AR technology. But can this technology be applied more broadly to support larger audiences in completing everyday tasks? Today, this question is being answered with a resounding "yes." A large variety of AR browser apps are already available on smartphones (e.g., Layar, Wikitudes, Junaio, and others). These apps are intended to deliver information related to *places of interest* in the user's environment, superimposed over the live video from the device's camera. The places of interest

Figure 1.20 AR browsers such as Yelp Monocle superimpose points of interest on a live video feed.

are either given in geo-coordinates and identified via the phone's sensors (GPS, compass readings) or identified by image recognition. AR browsers have obvious limitations, such as potentially poor GPS accuracy and augmentation capabilities only for individual points rather than full objects. Nevertheless, thanks to the proliferation of smartphones, these apps are universally available, and their use is growing, owing to the social networking capabilities built into the AR browsers. Figure 1.20 shows the AR browser *Yelp Monocle*, which is integrated into the social business review app *Yelp*.

Another compelling use case for AR browsing is simultaneous translation of foreign languages. This utility is now widely available in the *Google Translate* app (Figure 1.21). The user just has to select the target language and point the device camera toward the printed text; the translation then appears superimposed over the image.

Navigation

The idea of heads-up navigation, which does not distract the operator of a vehicle moving at high speeds from the environment ahead, was first considered in the context of military aircraft [Furness 1986]. A variety of see-through displays, which can be mounted to the visor of a pilot's helmet, have been developed since the 1970s. These devices, which are usually called heads-up displays, are mostly intended to show nonregistered information, such as the current speed or torque, but can also be used to show a form of AR. Military technology, however, is usually not directly applicable to the consumer market, which demands different ergonomics and pricing structures.

Figure 1.21 Google Translate superimposes spontaneous translations of text, recognized in real time, over the camera image.

With improved geo-information, it has become possible to overlay larger structures on in-car navigation systems, such as road networks. Figure 1.22 shows *Wikitude Drive*, a first-person car navigation system. The driving instructions are overlaid on top of the live video feed rather than being presented in a map-like view. The registration quality in this system is acceptable despite being based on smartphone sensors such as GPS, as the inertia of a car allows the system to predict the geography ahead with relative accuracy.

Figure 1.23 shows a parking assistant, which overlays a graphical visualization of the car trajectory onto the view of a rear-mounted camera.

Television

Many people likely first encountered AR as annotations to live camera footage brought to their homes via broadcast TV. The first and most prominent example of this concept is the virtual 1st & 10 line in American football, indicating the yardage needed for a first down, which is superimposed directly on the TV screencast of a game. While the idea and first patents for creating such on-field markers for football broadcasts date back to the late 1970s, it took until 1998 for the concept to be realized. The same concept of annotating TV footage with virtual overlays has successfully been applied to many other sports, including baseball, ice hockey, car racing,

Figure 1.22 Wikitude Drive superimposes a perspective view of the road ahead. Courtesy of Wikitude GmbH.

Figure 1.23 The parking assistant is a commercially available AR feature in many contemporary cars. Courtesy of Brigitte Ludwig.

Figure 1.24 Augmented TV broadcast of a soccer game. Courtesy of Teleclub and Vizrt, Switzerland (LiberoVision AG).

and sailing. Figure 1.24 shows a televised soccer game with augmentations. The audience in this incarnation of AR has no ability to vary the viewpoint individually. Given that the live action on the playing field is captured by tracked cameras, interactive viewpoint changes are still possible, albeit not under the end-viewer's control.

Several competing companies provide augmentation solutions for various broadcast events, creating convincing and informative live annotations. The annotation possibilities have long since moved beyond just sports information or simple line graphics, and now include sophisticated 3D graphics renderings of branding logos or product advertisements.

Using similar technology, it is possible—and, in fact, common in today's TV broadcasts—to present a moderator and other TV personalities in virtual studio settings. In this application, the moderator is filmed by tracked cameras in front of a green screen and inserted into a virtual rendering of the studio. The system even allows for interactive manipulation of virtual props.

Similar technologies are being used in the film industry, such as for providing a movie director and actors with live previews of what a film scene might look like after special effects or other compositing has been applied to the camera footage of a live set environment. This application of AR is sometimes referred to as Pre-Viz.

Figure 1.25 The lifestyle magazine *Red Bulletin* was the first print publication to feature dynamic content using AR. Courtesy of Daniel Wagner.

Advertising and Commerce

The ability of AR to instantaneously present arbitrary 3D views of a product to a potential buyer is already being welcomed in advertising and commerce. This technology can lead to truly interactive experiences for the customer. For example, customers in Lego stores can hold a toy box up to an AR kiosk, which then displays a 3D image of the assembled Lego model. Customers can turn the box to view the model from any vantage point.

An obvious target for AR is the augmentation of printed material, such as flyers or magazines. Readers of the *Harry Potter* novels know how pictures in the *Daily Prophet* newspaper come alive. This idea can be realized with AR by superimposing digital movies and animations on top of specific portions of a printed template. When the magazine is viewed on a computer or smartphone, the static pictures are replaced by animated sequences or movies (Figure 1.25).

AR can also be helpful for a sales person who is trying to demonstrate the virtues of a product (Figure 1.26). Especially for complex devices, it may be difficult to convey the internal operation with words alone. Letting a potential customer observe the animated interior allows for much more compelling presentations at trade shows and in show rooms alike.

Pictofit is a virtual dressing room application that lets users preview garments from online fashion stores on their own body (Figure 1.27). The garments are automatically adjusted to match

Figure 1.26 Marketing presentation of a Waeco air-conditioning service unit. Courtesy of magiclensapp.com.

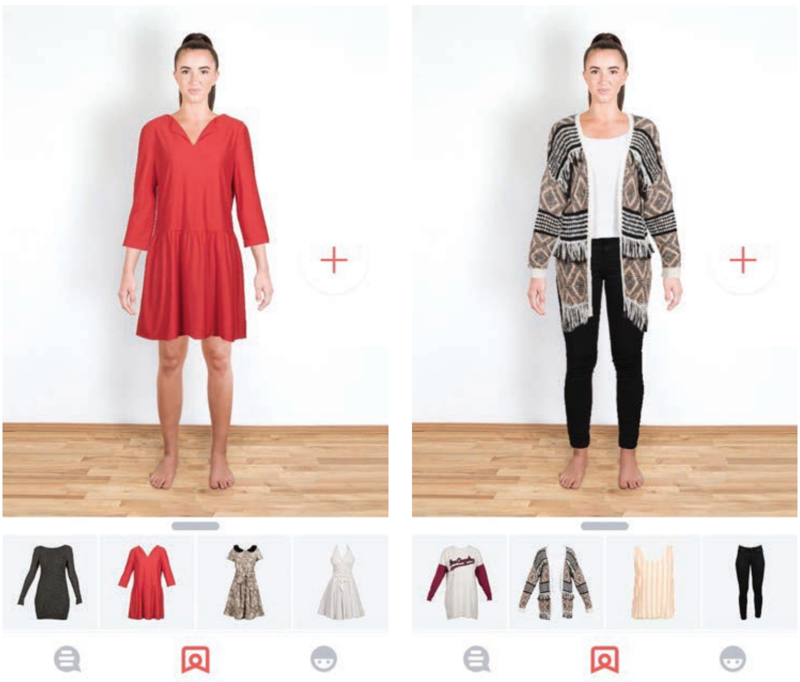

Figure 1.27 Pictofit can extract garment images from online shopping sites and render them to match an image of the customer. Courtesy of Stefan Hauswiesner, ReactiveReality.

the wearer's size. In addition, body measurements are estimated and made available to assist in the entry of purchase data.

Games

One of the first commercial AR games was *The Eye of Judgment*, an interactive trading card game for the Sony PlayStation 3. The game is delivered with an overhead camera, which picks up game cards and summons corresponding creatures to fight matches.

An important quality of traditional games is their tangible nature. Kids can turn their entire room into a playground, with pieces of furniture being converted into a landscape that supports physical activities such as jumping and hiding. In contrast, video games are usually confined to a purely virtual realm. AR can bring digital games together with the real environment. For example, Vuforia SmartTerrain (Figure 1.28) delivers a 3D scan of a real scene and turns it into a playing field for a "tower defense" game.

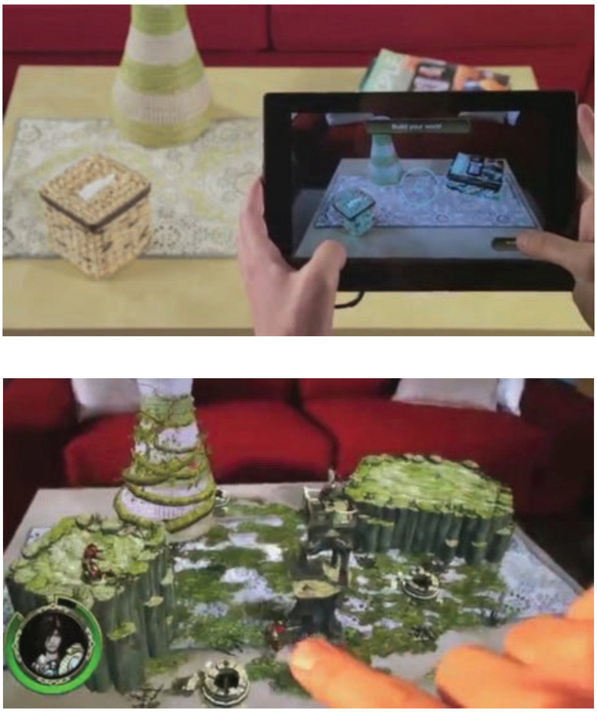

Figure 1.28 Vuforia SmartTerrain scans the environment and turns it into a game landscape.

Figure 1.29 Using a TV-plus-projector setup, the IllumiRoom extends the game world beyond the boundaries of the screen. Courtesy of Microsoft Research.

Microsoft's *IllumiRoom* [Jones et al. 2013] is a prototype of a projector-based AR game experience. It combines a regular TV set with a home-theater projector to extend the game world beyond the confines of the TV (Figure 1.29). The 3D game scene shown in the projection is registered with the one on the TV, but the projection covers a much wider field of view. While the player concentrates on the center screen, the peripheral field of view is also filled with dynamic images, leading to a greatly enhanced game experience.

Related Fields

In the previous section, we have highlighted a few AR applications. Other compelling examples of applications only tangentially match the definition we have given of AR. These applications often come from the related fields of mixed reality, ubiquitous computing, and virtual reality, which we briefly discuss here.

Mixed Reality Continuum

A user immersed in virtual reality experiences only virtual stimuli, for example, inside a CAVE (a room with walls consisting of stereoscopic back-projections) or when wearing a closed HMD. The space between reality and virtual reality, which allows real and virtual elements to be combined to varying degrees, is called **mixed reality**. In fact, some people prefer the term "mixed

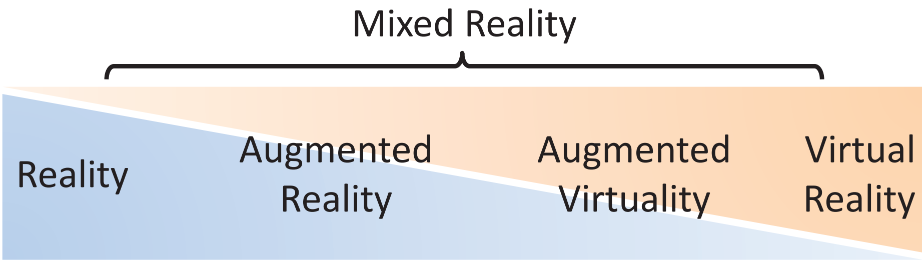

Figure 1.30 The mixed reality continuum captures all possible combinations of the real and virtual worlds.

reality" over "augmented reality," because they appreciate the broader and more encompassing notion of MR.

This view can be attributed to Milgram and Kishino [1994], who proposed a continuum (Figure 1.30) spanning from reality to virtual reality. They characterized MR as follows:

> [MR involves the] merging of real and virtual worlds somewhere along the "virtuality continuum" which connects completely real environments to completely virtual ones.

Benford et al. [1998] go one step further, arguing that a complex environment will often be composed of multiple displays and adjacent spaces, which constitute "mixed realities" (note the plural). These multiple spaces meet at "mixed reality boundaries."

According to this perspective, augmented reality contains primarily real elements and, therefore, is closer to reality. For example, a user with an AR app on a smartphone will continue perceiving the real world in the normal way, but with some additional elements presented on the smartphone. The real-world experience clearly dominates in such a case. The opposite concept, **augmented virtuality**, prevails when there are primarily virtual elements present. As an example, imagine an online role-playing game, where the avatars' faces are textured in real time with a video acquired from the player's face. Everything in this virtual game world, except the faces, is virtual.

Virtual Reality

At the far right end of the MR continuum, virtual reality immerses a user in a completely computer-generated environment. This removes any restrictions as to what a user can do or experience in VR. VR is now becoming increasingly popular for enhanced computer games. New designs for HMD gaming devices, such as the Oculus Rift or HTC Vive, are receiving a great deal of public attention. Such devices are also suitable for augmented virtuality applications. Consequently, AR and VR can easily coexist within the MR continuum. As we will see later, transitional interfaces can be designed to harness the combined advantages of both concepts.

Ubiquitous Computing

Mark Weiser proposed the concept of **ubiquitous computing** (ubicomp) in his seminal 1991 essay. His work anticipates the massive introduction of digital technology into everyday life. Contrasting ubicomp with virtual reality, he advocates bringing the "virtuality" of computer-readable data into the physical world via a variety of computer form factors, which should sound familiar to today's technology users: inch-scale "tabs," foot-scale "pads," and yard-scale "boards."

> Depending on the room, you may see more than 100 tabs, 10 or 20 pads, and one or two boards. This leads to our goal for initially deploying the hardware of embodied virtuality: hundreds of computers per room. [Weiser 1991]

This description includes the idea of mobile computing, which allows users to access digital information anytime and anywhere. However, it also predicts the "Internet of Things," in which all elements of our everyday environment are instrumented. Mackay [1998] has argued that augmented things should also be considered as a form of AR. Consider, for example, home automation, driver assistance systems in cars, and smart factories capable of mass customization. If such technology works well, it essentially disappears from our perception. The first two sentences of Weiser's 1991 article succinctly express this model:

> The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.

Ubicomp is primarily intended as "calm computing"; that is, human attention or control is neither required nor intended. However, at some point, control will still be necessary. A human operator away from a desktop computer, for example, may need to steer complex equipment. In such a situation, an AR interface can directly present status updates, telemetry information, and control widgets in a view of the real environment. In this sense, AR and ubicomp fit extremely well: AR is the ideal *user interface* for ubicomp systems.

According to Weiser, VR is the opposite of ubicomp. Weiser notes the monolithic nature of VR environments, such as a CAVE, which isolate a user from the real world. However, Newman et al. [2007] suggest that ubicomp actually combines two important characteristics: **virtuality** and **ubiquity**. Virtuality, as described by the MR continuum, expresses the degree to which virtual and reality are mixed. Weiser considers location and place as computational inputs. Thus, ubiquity describes the degree to which information access is independent from being in a fixed place (a terminal). Based on these understandings, we can arrange a family of technologies in a "Milgram–Weiser" chart as shown in Figure 1.31.

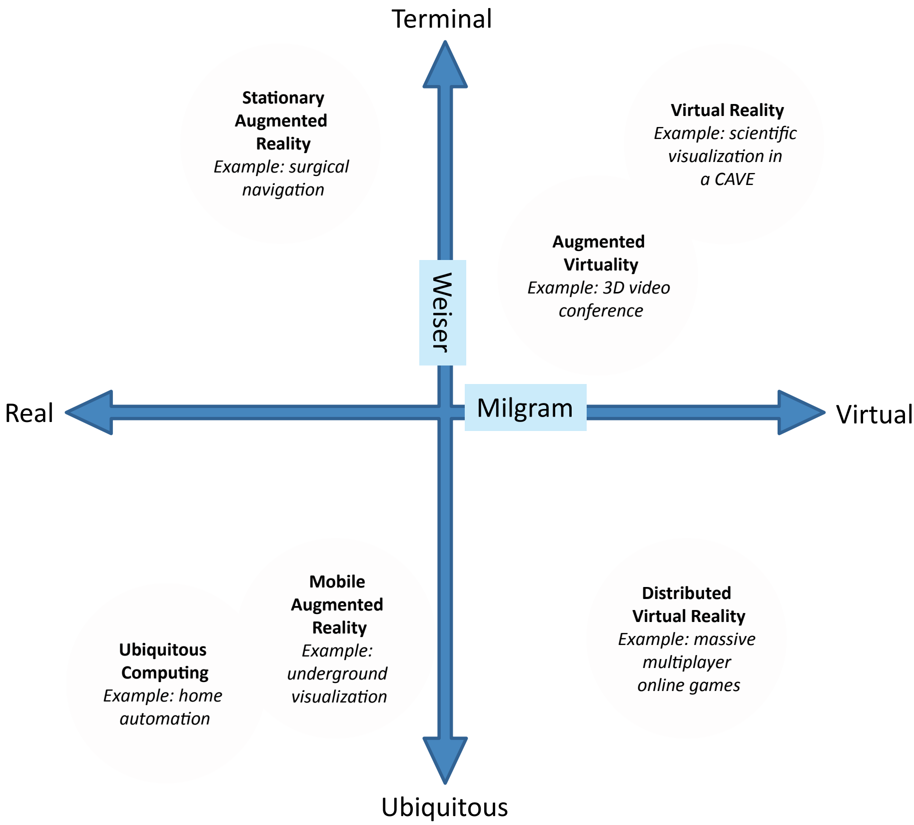

Figure 1.31 The Milgram–Weiser chart visualizes the relationships of various user interface paradigms.

Summary

In this chapter, we provided an introduction to the research field and practical occurrences of augmented reality. For a working definition, augmented reality relies on three key components: (1) the combination of virtual and real information, with the real world as the primary place of action and (2) interactive, real-time updates of (3) virtual information registered in 3D with the physical environment. Different technologies can be used to realize such a concept. The first part of this book provides an overview of technologies for displays (Chapter 2), tracking technologies (Chapters 3, 4, and 5) and graphics (Chapters 6 and 7). The second part of the book (Chapters 8 through 14) deals with interactive techniques.

We also presented a brief history of the field and then went on a whirlwind tour of AR application examples, with the goal of suggesting the enormous potential that AR holds as an interface metaphor to computing in the physical world (sometimes referred to as situated computing). While many specific application possibilities exist, such as AR for equipment maintenance or AR for surgery, one can also envision AR turning into a more general interface paradigm, redefining the overall browsing experience for computing in the physical world. Application examples from the domains of personal information display and navigation hint at that potential.

We concluded this chapter with a discussion of related fields. In doing so, we placed AR within the scope of Milgram's mixed reality continuum and contrasted AR with Weiser's concept of ubiquitous computing.

INDEX

B

C

G

H

I

J

K

L

M

N

O

P

Q

R

S

U

V